

Automated Verification, Synthesis and Correction of Concurrent Systems via MSO Logic

Mateus de Oliveira Oliveira

KTH Royal Institute of Technology
mdeoliv@kth.se

Abstract. In this work we provide algorithmic solutions to five fundamental problems concerning the verification, synthesis and correction of concurrent systems that can be modeled by bounded p/t -nets. We express concurrency via partial orders and assume that behavioral specifications are given via monadic second order logic. A c -partial-order is a partial order whose Hasse diagram can be covered by c paths. For a finite set T of transitions, we let $\mathcal{P}(c, T, \varphi)$ denote the set of all T -labelled c -partial-orders satisfying φ . If $N = (P, T)$ is a p/t -net we let $\mathcal{P}(N, c)$ denote the set of all c -partially-ordered runs of N . A (b, r) -bounded p/t -net is a b -bounded p/t -net in which each place appears repeated at most r times. We solve the following problems:

1. **Verification:** given an MSO formula φ and a bounded p/t -net N determine whether $\mathcal{P}(N, c) \subseteq \mathcal{P}(c, T, \varphi)$, whether $\mathcal{P}(c, T, \varphi) \subseteq \mathcal{P}(N, c)$, or whether $\mathcal{P}(N, c) \cap \mathcal{P}(c, T, \varphi) = \emptyset$.
2. **Synthesis from MSO Specifications:** given an MSO formula φ , synthesize a semantically minimal (b, r) -bounded p/t -net N satisfying $\mathcal{P}(c, T, \varphi) \subseteq \mathcal{P}(N, c)$.
3. **Semantically Safest Subsystem:** given an MSO formula φ defining a set of safe partial orders, and a b -bounded p/t -net N , possibly containing unsafe behaviors, synthesize the safest (b, r) -bounded p/t -net N' whose behavior lies in between $\mathcal{P}(N, c) \cap \mathcal{P}(c, T, \varphi)$ and $\mathcal{P}(N, c)$.
4. **Behavioral Repair:** given two MSO formulas φ and ψ , and a b -bounded p/t -net N , synthesize a semantically minimal (b, r) -bounded p/t net N' whose behavior lies in between $\mathcal{P}(N, c) \cap \mathcal{P}(c, T, \varphi)$ and $\mathcal{P}(c, T, \psi)$.
5. **Synthesis from Contracts:** given an MSO formula φ^{yes} specifying a set of good behaviors and an MSO formula φ^{no} specifying a set of bad behaviors, synthesize a semantically minimal (b, r) -bounded p/t -net N such that $\mathcal{P}(c, T, \varphi^{yes}) \subseteq \mathcal{P}(N, c)$ but $\mathcal{P}(c, T, \varphi^{no}) \cap \mathcal{P}(N, c) = \emptyset$.

Key words: System Synthesis, Verification of Concurrent Systems, Automated Repair, Monadic Second Order Logic, Partial Orders, Slice Theory

1 Introduction

Model checking and *system synthesis* are two complementary paradigms that are widely used to provide correctness guarantees for computational systems. On the one hand, the goal of model checking is to verify whether the behavior of a given system is in accordance with a given specification [10,11,31,33]. On the other hand, the goal of system synthesis is to mechanically construct a system from a behavioral specification [8,18,26,28,32]. When combined, model checking and synthesis can be used as primitives for the development of powerful methodologies aimed at the mechanical correction of bugs, such as system repair [23,24,34,37]. In this work we develop a combined

theory of *model checking* and *system synthesis* that is fully compatible with the partial order theory of concurrency. Our systems are modeled via bounded place/transition nets, while our behavioral specifications are given in monadic second order logic. We solve five fundamental problems lying in the intersection of system verification, system synthesis and system repair. First we show how to compare the partial order behavior of bounded p/t -nets with partial order behaviors specified via MSO formulas. Second, we show how to synthesize bounded p/t -nets from MSO-definable sets of partial orders. Third, we show how to obtain the *semantically safest subsystem* of a bounded p/t -net with respect to a MSO specification. Fourth, we transpose the methodology of program repair introduced by Jobstmann and von Essen [37] to the context of bounded p/t -nets with partial-order runs. Finally, we show how to synthesize bounded p/t -nets from partial-order contracts. Before giving a precise definition of each of the problems described above, we briefly introduce the main elements of our model.

Bounded place/transition nets: Petri nets [30], also known as place/transition-nets are recognized as an elegant mathematical formalism for the specification of concurrent systems. During the last four decades, p/t -nets have found applications in the modeling of real time fault tolerant systems, faulty critical systems, communication protocols, logic controllers, and many others types of computing systems [29,38]. A p/t -net consists of a multiset of places, which are initially loaded with a set of tokens, and a set of transitions. The disposition of the tokens among the places of a p/t -net determine which transitions are allowed to fire. A transition, by its turn, when fired, removes tokens from some places and adds tokens to some places. In this work we will be concerned with the partial order theory of bounded p/t -nets. We say that a p/t -net is b -bounded if after firing any sequence of transitions, the number of tokens in each of its places remains bounded by b , and that a p/t -net is (b, r) -bounded if it is b -bounded and if each place occurs in it at most r times. We will define p/t -nets more precisely in Section 2.

Partial Orders: When concurrency is interpreted accordingly to the interleaving semantics, the execution of concurrent actions is identified with the non-deterministic choice among all possible orders in which such actions can occur. Although satisfactory for many applications of practical relevance, this point of view has some drawbacks. First, the interleaving semantics is not compatible with the notion of action refinement, in which an atomic action is replaced by a set of sub-actions [35]. Second, this point of view is not appropriate to model concurrent scenarios in which several users have concurrent read/write access to databases [19] nor to model the behavior of read/write operations in multiprocessors that implement *weak memory models* [1].

A well established point of view which is able to overcome these drawbacks is to represent both concurrency and causality as partially ordered sets of events [21,20,36]. The partial order imposed on a set of events can be interpreted according to two standard semantics. According to the first, the *causal semantics*, an event v is smaller than an event v' if v' causally depends on the occurrence of v . According to the second, the *execution semantics*, the fact that v is smaller than v' simply indicates that v does not occur after v' . However, in this case the event v may not necessarily be one of the causes of the event v' . In this work, we will study the partial order behavior of bounded p/t -nets according to both semantics.

***c*-Partial-Orders:** We introduce a new parameterization for the study of the partial order behavior of concurrent systems. Recall that the Hasse diagram of a partial order ℓ is the directed acyclic graph H with the least number of edges whose transitive closure equals ℓ . We say that a partial order ℓ is a *c*-partial-order if its Hasse diagram H can be covered by c paths. In other words if there exist paths p_1, \dots, p_k in H such that $H = \bigcup_{i=1}^c p_i$. We notice that the paths are not assumed to be edge disjoint nor vertex disjoint. We let $\mathcal{P}_{cau}(N, c)$ denote the set of all *c*-partial-orders that can be associated to a *p/t*-net N according to the causal semantics, and by $\mathcal{P}_{ex}(N, c)$ the set of all *c*-partial-orders that can be associated with N according to the execution semantics. Intuitively, the parameter c characterizes the thickness of the partial order, and provides a width measure that is stronger and more algorithmically friendly than the traditional notion of width used in partial order theory. We observe that the execution behavior $\mathcal{P}_{ex}(N, 1)$ is simply the set of all possible firing sequences of N . If N is a b -bounded *p/t*-net with n places then the set $\mathcal{P}_{cau}(N, b \cdot n)$ already comprises all possible causal runs of N . We contrast this observation with the fact that there are very simple examples¹ of *p/t*-nets whose execution behavior $\mathcal{P}_{ex}(N, c)$ is strictly contained into $\mathcal{P}_{ex}(N, c + 1)$ for each $c \in \mathbb{N}$.

Monadic Second Order Logic of Graphs: The monadic second order logic over partial orders extends first order logic by adding the possibility of quantifying over sets of vertices. The role of MSO logic in the study of the partial order behavior of concurrent systems was emphasized in [27] in the context of the theory of message sequence graphs. Let T be a finite set of symbols, which should be regarded as labels of transitions in a concurrent system. We say that a partial order ℓ is a T -labeled partial order if each of its nodes are labeled with some element of T . Let φ be an MSO formula expressing a property of T -labeled partial orders, and let $c \in \mathbb{N}$. We denote by $\mathcal{P}(c, T, \varphi)$ the set of all T -labeled *c*-partial-orders satisfying φ . In this work the connection between the *c*-partial-order behavior of bounded *p/t*-nets and MSO logic will be established via a formalism called slice automaton (Section 3). In the context of this paper, slice automata should be regarded as a generalization of message sequence graphs which is suitable for the representation of the *c*-partial-order behavior of bounded *p/t*-nets. Indeed, we will show that for each MSO formula φ , the set of all *c*-partial-orders satisfying φ can be represented by a slice automaton. The connection with bounded *p/t*-nets stems from a result previously proved by us [17] stating that the *c*-partial-order behavior of bounded *p/t*-nets can also be effectively represented via slice automata.

In the next five subsections we will state our main results and establish further connections with existing literature.

1.1 Verification of the Partial Order Behavior of Bounded *p/t*-Nets.

Suppose we have a concurrent system modeled by a b -bounded *p/t*-net N and let φ be an MSO formula. In Theorem 1.1 below we address three verification results. First, assuming that φ defines a set of faulty behaviors, we can mechanically determine whether or not some of the partial order runs of N is faulty (Theorem 1.1.i). Second, on the contrapositive, assuming that φ specifies a set of good partial order behaviors, we can

¹ For instance a net consisting of two places p_1, p_2 , initialize with a unique token each, and two transitions t_1, t_2 such that t_i takes one token from p_i and puts it back on p_i .

test whether or not all behaviors of N are good (Theorem 1.1.ii). Third, assuming that φ specify a set of desired partial order behaviors, we can decide whether the partial order behavior of N comprises all partial orders specified by φ (Theorem 1.1.iii). All three verification results hold with respect with both the execution and the causal semantics. Below, we let the variable sem be equal to ex if we are considering the execution semantics and equal to cau if we are considering the causal semantics. A precise definition of how partial orders are assigned to p/t -nets according to each of these semantics will be given in Section 2.

Theorem 1.1 (Verification). *Let φ be an MSO formula, N be a b -bounded p/t -net, $c \in \mathbb{N}$, and $sem \in \{ex, cau\}$.*

- i) *One may effectively determine whether $\mathcal{P}_{sem}(N, c) \cap \mathcal{P}(c, T, \varphi) = \emptyset$.*
- ii) *One may effectively determine whether $\mathcal{P}_{sem}(N, c) \subseteq \mathcal{P}(c, T, \varphi)$.*
- iii) *One may effectively determine whether $\mathcal{P}(c, T, \varphi) \subseteq \mathcal{P}_{sem}(N, c)$.*

Notice that Theorems 1.1.i and 1.1.ii can be reduced to each other since $\mathcal{P}_{sem}(N, c) \cap \mathcal{P}(c, T, \varphi) = \emptyset$ if and only if $\mathcal{P}_{sem}(N, c) \subseteq \mathcal{P}(c, T, \neg\varphi)$. Theorem 1.1 addresses for the first time safety and conformance tests of the both the execution and the causal behaviors of general bounded p/t -nets. We notice that in the special case of *pure*² bounded p/t -nets, the model checking of the causal behavior was addressed in [2] using the machinery of vector addition systems with states. In our notation, this corresponds to testing whether $\mathcal{P}_{cau}(N, n \cdot b) \cap \mathcal{P}(n \cdot b, T, \varphi) = \emptyset$, where n is the number of places of N and b its bound. However, as pointed out in [3], pure p/t -nets are a rather restricted subclass of p/t -nets, since they are not able to model for instance, waiting loops in communication protocols. The results in [2] are not able to address the model checking of p/t -nets according to the execution semantics, nor to provide an analog of Theorem 1.1.iii with respect to neither the causal nor the execution semantics.

1.2 Synthesis of Bounded p/t -Nets from MSO Specifications

In our second result (Theorem 1.2) we address the synthesis of p/t -nets from MSO definable sets of c -partial-orders. We say that a p/t -net is (b, r) -bounded if each place occurs with multiplicity at most r and if each place has at most b -tokens on each legal marking of N . Let $sem \in \{ex, cau\}$. A (b, r) -bounded p/t -net N is c -*sem*-minimal for a partial order language \mathcal{P} , if $\mathcal{P} \subseteq \mathcal{P}_{sem}(N, c)$ and if there is no other (b, r) -bounded p/t -net N' with $\mathcal{P} \subseteq \mathcal{P}_{sem}(N', c) \subsetneq \mathcal{P}_{sem}(N, c)$.

Theorem 1.2 (Synthesis). *Let φ be an MSO formula, T be a finite set of transitions, and $b, c, r \in \mathbb{N}$. Then one can effectively determine whether there exists a (b, r) -bounded p/t -net N which is c -sem-minimal for $\mathcal{P}(c, T, \varphi)$. In the case such a net N exists one can effectively construct it.*

Theorem 1.2 addresses for the first time the synthesis of bounded p/t -nets from MSO specifications. Observe that the minimality condition imposed in theorem 1.2 implies that in the case that it is not possible to synthesize a net precisely matching the specification φ , it is still possible to synthesize a net with the fewest number of bad partial-order runs as possible. We observe that the parameter r in Theorem 1.2 is only relevant when

² A p/t -net is pure if no transition takes a token from a place and puts it in the same place.

considering the causal semantics. Indeed, adding repeated places to p/t -nets does not change their execution behavior. However the addition of repeated places can indeed increase the causal behavior of p/t -net [15]. We also observe that when considering the synthesis with the execution semantics all c -ex-minimal p/t -nets for a partial order language \mathcal{P} have the same partial order behavior. However when considering the causal semantics, there may exist two p/t -nets N_1 and N_2 whose behavior is c -cau-minimal for \mathcal{P} , but for which $\mathcal{P}(N_1, c) \neq \mathcal{P}(N_2, c)$.

It is worth comparing Theorem 1.2 with existing literature. When considering the interleaving semantics of bounded p/t -nets, the synthesis problem from regular languages was studied extensively in [3,4,13,14] via a set of combinatorial techniques called *theory of regions*. Thus, via Buchi-Elgot Theorem stating that MSO Logic over strings is as expressive as regular languages [7], the theory of regions can be used to synthesize nets whose interleaving behavior satisfies a given MSO formula over *strings*. In our notation this corresponds to synthesizing a bounded net N whose 1-execution-behavior $\mathcal{P}_{ex}(N, 1)$ is 1-ex-minimal with respect to $\mathcal{P}(\varphi, T, 1)$. Here we solve the synthesis problem from MSO languages for any $c \geq 1$, and with respect to both the causal and execution semantics. It is worth noting that the synthesis of bounded p/t -nets with the execution semantics from certain restricted partial order formalisms that are not able to represent the behavior of bounded p/t -nets was considered in [5,6], but no connection with logic was established therein. The synthesis of bounded p/t -nets from a mathematical object that is able to fully represent the causal behavior of bounded p/t -nets was solved by us in [15], solving in this way an open problem stated in [25]. This mathematical object is called slice automaton, and will be described in Section 3. The proof of Theorem 1.2 follows by establishing a non-trivial connection between monadic second order logic and slice automata.

1.3 Semantically Safest Subsystem

Suppose that we have in hands an MSO formula φ specifying a set of safe behaviors, and a concurrent system specified by a (b, r) -bounded p/t -net N . Suppose that after verifying N according to Theorem 1.1 we discover that some runs of N are faulty, i.e., do not satisfy φ . What should we do? Discard N , and try to re-project a new system from scratch? In the next theorem (Theorem 1.3) we will show that we may still be able to save N by automatically synthesizing the best (b, r) -bounded p/t -net N' whose partial order behavior lies in between $\mathcal{P}_{sem}(N, c) \cap \mathcal{P}(\varphi, c)$ and $\mathcal{P}_{sem}(N, c)$. In other words, the partial order behavior of N' is a subset of the partial order behavior of N which preserves all safe runs of N . Additionally, the partial order behavior of N' has as few unsafe partial-order runs as possible. We call N' the *semantically safest subsystem* of N . We notice that the net N' does not need to be a sub-net of N , and indeed N' can have even more places than N . Only the behavior of N' is guaranteed to be a subset of the behavior of N .

Theorem 1.3 (Semantically Safest Subsystem). *Let $c, b, r \in \mathbb{N}$ and $sem \in \{ex, cau\}$. Given a (b, r) -bounded p/t -net $N = (P, T)$ and an MSO formula φ , we may automatically synthesize a (b, r) -bounded p/t -net N' such that*

- i) N' is c -sem-minimal for $\mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c)$,
- ii) $\mathcal{P}_{sem}(N', c) \subseteq \mathcal{P}_{sem}(N, c)$.

We consider that our notion of *semantically safest subsystem* is appropriate for three reasons. First, as mentioned above, 1.3.i and 1.3.ii imply that $\mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c) \subseteq \mathcal{P}(N', c) \subseteq \mathcal{P}(N, c)$. Second, the minimality condition says that if there is a (b, r) -bounded p/t -net N' whose c -partial-order behavior precisely matches $\mathcal{P}_{sem}(N, c) \cap \mathcal{P}(\varphi, c)$ then such a p/t -net will be returned. In this case, our synthesis algorithm completely corrects the original p/t -net. Finally, but not less important, if all c -partially-ordered runs of N indeed satisfy φ , then our synthesis algorithm returns a net N' satisfying $\mathcal{P}_{sem}(N', c) = \mathcal{P}_{sem}(N, c)$. Thus the set of c -partial order behaviors of the synthesized net does not change if the original net is already correct (although the structure of the net per se may change). In Subsection 1.4 below we consider a related problem that finds analogies with the field of automatic program repair.

1.4 Behavioral Repair

During the last decade a substantial amount of effort has been devoted to the development of methodologies for the automatic correction of bugs in computational systems [9,23,24,34]. Very recently, in the context of reactive systems, Jobstmann and von Esen have combined system synthesis and model checking to develop a methodology of program repair that preserves semantically correct runs [37]. Within their methodology, given two LTL formulas φ and ψ and a reactive system S , one is asked to automatically synthesize a system S' whose behavior is lower bounded by $\mathcal{L}(\varphi) \cap \mathcal{L}(S)$ and upper bounded by $\mathcal{L}(\psi)$. Intuitively, while φ specifies a set of correct behaviors that should be preserved whenever present in the original system, the formula ψ specifies the set of behaviors that are allowed to be present in the repaired system. In Theorem 1.4 below we transpose the semantically preserving repair methodology devised in [37] to the realm of bounded p/t -nets with the partial order semantics.

Theorem 1.4 (Behavioral Repair). *Let $c, b, r \in \mathbb{N}$ and $sem \in \{ex, cau\}$. Given a (b, r) -bounded p/t -net $N = (P, T)$ and an MSO formula φ , we may automatically determine whether there exists a (b, r) -bounded p/t -net N' such that*

- i) N' is c -sem-minimal for $\mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c)$,
- ii) $\mathcal{P}_{sem}(N', c) \subseteq \mathcal{P}_{sem}(c, T, \psi)$.

In the case such a net exists, one may automatically construct it.

While 1.4.i and 1.4.ii imply that $\mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c) \subseteq \mathcal{P}(N', c) \subseteq \mathcal{P}(c, T, \psi)$, the minimality condition in 1.4.i implies that if N' is successfully synthesized, then its behavior has as few partial-order runs contradicting φ as possible.

1.5 Synthesis from Partial Order Contracts

Suppose that we are in the early stages of development of a concurrent system. We have arrived to the conclusion that every behavior satisfying a given MSO formula φ^{yes} should be present in the system, but that no behavior in the system should satisfy a formula φ^{no} . Clearly we require that $\mathcal{P}(\varphi^{yes}) \cap \mathcal{P}(\varphi^{no}) = \emptyset$. We say that the pair $(\varphi^{yes}, \varphi^{no})$ is a partial order contract. We can try to develop a first prototype of our system by automatically synthesizing a (b, r) -bounded p/t -net N containing all c -partial orders specified by φ^{yes} but no partial order in φ^{no} . The next theorem says that if such a net exists, then it can be automatically constructed.

Theorem 1.5 (Synthesis from Contracts). *Let φ^{yes} and φ^{no} be MSO formulas with $\mathcal{P}(c, T, \varphi^{yes}) \cap \mathcal{P}(c, T, \varphi^{no}) = \emptyset$. Then one may automatically determine whether there exists a (b, r) -bounded p/t-net N such that $\mathcal{P}(c, T, \varphi^{yes}) \subseteq \mathcal{P}_{sem}(N, c)$ and $\mathcal{P}(c, T, \varphi^{no}) \cap \mathcal{P}_{sem}(N, c) = \emptyset$. In the case such a net exists one may construct it.*

2 p/t-Nets and their Partial Order Semantics

Let T be a finite set of transitions. Then a place over T is a triple $p = (p_0, \check{p}, \hat{p})$ where p_0 denotes the initial number of tokens in p and $\check{p}, \hat{p} : T \rightarrow \mathbb{N}$ are functions which denote the number of tokens that each transition $t \in T$ respectively puts in and takes from p . A p/t-net over T is a pair $N = (P, T)$ where T is a set of transitions and P a finite multi-set of places over T . We assume through this paper that for each transition $t \in T$, there exist places $p_1, p_2 \in P$ for which $\check{p}_1(t) > 0$ and $\hat{p}_2(t) > 0$. A marking of N is a function $m : P \rightarrow \mathbb{N}$. A transition t is enabled at marking m if $m(p) \geq \hat{p}(t)$ for each $p \in P$. The occurrence of an enabled transition at marking m gives rise to a new marking m' defined as $m'(p) = m(p) - \hat{p}(t) + \check{p}(t)$. The initial marking m_0 of N is given by $m_0(p) = p_0$ for each $p \in P$. A sequence of transitions $t_1 t_2 \dots t_n$ is an occurrence sequence of N if there exists a sequence of markings $m_0 m_1 \dots m_n$ such that for each $i \in \{1, \dots, n\}$, t_i is enabled at m_{i-1} and if m_i is obtained by the firing of t_i at marking m_{i-1} . A marking m is legal if it is obtained from m_0 by the firing of an occurrence sequence of N . A place p of N is b -bounded if $m(p) \leq b$ for each legal marking m of N . A net N is b -bounded if each of its places is b -bounded. The union of two p/t-nets $N_1 = (P_1, T)$ and $N_2 = (P_2, T)$ having a common set of transitions T is the p/t-net $N_1 \cup N_2 = (P_1 \cup P_2, T)$. Observe that since we are dealing with the union of multisets, if a place p occurs with multiplicity r_1 in P_1 and with multiplicity r_2 in P_2 then the same place will occur with multiplicity $r_1 + r_2$ in $P_1 \cup P_2$.

The notion *process*, upon which the partial order semantics of p/t-nets is derived, is defined in terms of objects called *occurrence nets*. An occurrence net is a DAG $O = (B \dot{\cup} V, F)$ where the vertex set $B \dot{\cup} V$ is partitioned into a set B , whose elements are called conditions, and a set V , whose elements are called events. The edge set $F \subseteq (B \times V) \cup (V \times B)$ is restricted in such a way that for every condition $b \in B$,

$$|\{(b, v) \mid v \in V\}| \leq 1 \quad \text{and} \quad |\{(v, b) \mid v \in V\}| \leq 1.$$

In other words, conditions in an occurrence net are unbranched. For a condition $b \in B$ we let $InDegree(b)$ denote the number of edges having b as target. A process of a p/t-net is an occurrence net in which conditions are labeled with places of N and events are labeled with transitions of N in such a way that the number of conditions labeled by a place $p \in N$ which immediately precede (follows) an event labeled by a transition t is equal to $\hat{p}(t)$ ($\check{p}(p)$). We define processes more precisely below.

Definition 2.1 (Process [22]). *A process of a p/t-net $N = (P, T)$ is a labeled DAG $\pi = (B \dot{\cup} V, F, \rho)$ where $(B \dot{\cup} V, F)$ is an occurrence net and $\rho : (B \cup V) \rightarrow (P \cup T)$ is a labeling function satisfying the following properties.*

1. *Places label conditions and transitions label events.*

$$\rho(B) \subseteq P \quad \rho(V) \subseteq T$$

2. For every $v \in V$, and every $p \in P$,

$$|\{(b, v) \in F : \rho(b) = p\}| = \hat{p}(\rho(v)) \quad \text{and} \quad |\{(v, b) \in F : \rho(b) = p\}| = \check{p}(\rho(v))$$

3. For every $p \in P$,

$$|\{b | \text{InDegree}(b) = 0, \rho(b) = p\}| = p_0.$$

Let $R \subseteq X \times X$ be a binary relation on a set X . We denote by $\text{tc}(X)$ the transitive closure of R . If $\pi = (B \cup V, F, \rho)$ is a process then the *causal order* of π is the partial order $\ell_\pi = (V, \text{tc}(F)|_{V \times V}, \rho|_V)$ which is obtained by taking the transitive closure of F and subsequently by restricting $\text{tc}(F)$ to pairs of events of V . In other words the causal order of a process π is the partial order induced by π on its events. We denote by $\mathcal{P}_{\text{cau}}(N)$ the set of all partial orders derived from processes of N . We say that $\mathcal{P}_{\text{cau}}(N)$ is the causal language of N .

$$\mathcal{P}_{\text{cau}}(N) = \{\ell_\pi | \pi \text{ is a process of } N\}$$

Observe that several processes of a p/t -net N may correspond to the same partial order. A sequentialization of a partial order ℓ is any partial order $\ell' = (V, <', l)$ for which $< \subseteq <'$. If N is a p/t -net then an *execution* of N is any sequentialization of a causal order in $\mathcal{P}_{\text{cau}}(N)$. We denote by $\mathcal{P}_{\text{ex}}(N)$ the set of all executions of N .

$$\mathcal{P}_{\text{ex}}(N) = \{\ell | \ell \text{ is a sequentialization of a causal order in } \mathcal{P}_{\text{cau}}(N)\}.$$

We denote by $\mathcal{P}_{\text{ex}}(N, c)$ the set of all c -partial orders in $\mathcal{P}_{\text{ex}}(N)$ and by $\mathcal{P}_{\text{cau}}(N, c)$ the set of all c -partial orders in $\mathcal{P}_{\text{cau}}(N)$. We notice that when considering the execution semantics of a b -bounded p/t -net $N = (P, T)$, the set $\mathcal{P}_{\text{ex}}(N, 1)$ is simply the set of all occurrence sequences of N . Additionally, the inclusion $\mathcal{P}_{\text{ex}}(N, c) \subseteq \mathcal{P}_{\text{ex}}(N, c + 1)$ may be proper for infinitely many values of c . In other words the execution behavior of a p/t -net may increase infinitely often with an increase in the parameter c . On the other hand, when considering the causal semantics of N , it can be shown [15] that

$$\mathcal{P}_{\text{cau}}(N, b \cdot |P|) = \mathcal{P}_{\text{cau}}(N, b \cdot |P| + i) = \mathcal{P}_{\text{cau}}(N)$$

for any $i \in \mathbb{N}$. Thus the causal behavior of a b -bounded p/t -net stabilizes for $c = b \cdot |P|$.

3 Regular Slice Languages

A slice $\mathbf{S} = (V, E, l, s, t)$ is a DAG³ where $V = I \dot{\cup} C \dot{\cup} O$ is a set of vertices partitioned into an in-frontier I , a center C and an out-frontier O ; E is a set of edges, $s, t : E \rightarrow V$ are functions that associate to each edge $e \in E$ a source vertex e^s and a target vertex e^t , and $l : V \rightarrow T \cup \mathbb{N}$ is a function that labels the center vertices in C with elements of a finite set T , and the in- and out-frontier vertices with positive integers in such a

³ A generalization of slices to arbitrary digraphs was considered in [17], but in this work we are only interested in slices that give rise to DAGs.

way that $l(I) = \{1, \dots, |I|\}$ and $l(O) = \{1, \dots, |O|\}$. Additionally, we require that each frontier-vertex v in $I \cup O$ is the endpoint of exactly one edge $e \in E$ and that no edge has both endpoints in the same frontier. Finally, in this work, we consider that the edges are directed from the in-frontier to the out frontier. In other words, for each edge $e \in E$, $e^s \in I \cup C$ and $e^t \in C \cup O$. From now on we will omit the source and target functions s and t from the specification of a slice and write simply $\mathbf{S} = (V, E, l)$.

A slice $\mathbf{S}_1 = (V_1, E_1, l_1)$ with frontiers (I_1, O_1) can be glued to a slice $\mathbf{S}_2 = (V_2, E_2, l_2)$ with frontiers (I_2, O_2) provided $|O_1| = |I_2|$. In this case the glueing gives rise to the slice $\mathbf{S}_1 \circ \mathbf{S}_2$ with frontiers (I_1, O_2) which is obtained by fusing, for each $i \in \{1, \dots, |O_1|\}$, the unique edge $e_1 \in E_1$ for which $l_1(e_1^t) = i$ with the unique edge $e_2 \in E_2$ for which $l_2(e_2^s) = i$. Formally, the fusion of e_1 with e_2 proceeds as follows. First we create an edge e_{12} . Then we set $e_{12}^s = e_1^s$ and $e_{12}^t = e_2^t$. Finally we delete both e_1 and e_2 . Thus in the glueing process the vertices in the glued frontiers disappear.

A *unit slice* is a slice with exactly one vertex in its center. A slice is *initial* if it has empty in-frontier and *final* if it has empty out-frontier. The width of a slice \mathbf{S} with frontiers (I, O) is defined as $w(\mathbf{S}) = \max\{|I|, |O|\}$. If T is a finite alphabet of symbols, then we let $\vec{\Sigma}(c, T)$ be the set of all unit slices of width at most c whose unique center vertex is labeled with an element of T . Observe that $\vec{\Sigma}(c, T)$ is finite and has asymptotically $|T| \cdot 2^{O(c \log c)}$ slices. We let $\vec{\Sigma}(c, T)^*$ denote the free monoid generated by $\vec{\Sigma}(c, T)$. We should *emphasize* that at this point the operation of the free monoid is simply the concatenation \mathbf{SS}' of slices and *should not be confused* with the composition $\mathbf{S} \circ \mathbf{S}'$. Thus the elements of $\vec{\Sigma}(c, T)^*$ are simply sequences $\mathbf{S}_1 \mathbf{S}_2 \dots \mathbf{S}_n$ of slices regarded as dumb letters. Additionally, the identity element of this monoid is simply the empty string λ , for which $\lambda \mathbf{S} = \mathbf{S} = \mathbf{S} \lambda$.

We let $\mathcal{L}(\vec{\Sigma}(c, T))$ be the set of all sequences of slices $\mathbf{S}_1 \mathbf{S}_2 \dots \mathbf{S}_n \in \vec{\Sigma}(c, T)^*$ for which \mathbf{S}_i can be composed with \mathbf{S}_{i+1} for $i \in \{1, \dots, n-1\}$, and for which \mathbf{S}_1 is initial and \mathbf{S}_n is final. We call the elements of $\mathcal{L}(\vec{\Sigma}(c, T))$ *unit decompositions*. If $\mathbf{U} = \mathbf{S}_1 \mathbf{S}_2 \dots \mathbf{S}_n$ is a unit decomposition in $\mathcal{L}(\vec{\Sigma}(c, T))$ then we denote by $\mathring{\mathbf{U}} = \mathbf{S}_1 \circ \mathbf{S}_2 \circ \dots \circ \mathbf{S}_n$ the DAG obtained from \mathbf{U} by composing all of its slices. A *slice language* over $\vec{\Sigma}(c, T)$ is any subset of $\mathcal{L}(\vec{\Sigma}(c, T))$. The width $w(\mathbf{U})$ of a unit decomposition $\mathbf{U} = \mathbf{S}_1 \mathbf{S}_2 \dots \mathbf{S}_n$ is the maximum width of a slice occurring in \mathbf{U} : $w(\mathbf{U}) = \max_i w(\mathbf{S}_i)$. Each slice language \mathcal{L} represents a possibly infinite family of DAGs $\mathcal{L}_{\mathcal{G}}$ which is obtained by composing the slices in each unit decomposition in \mathcal{L} .

$$\mathcal{L}_{\mathcal{G}} = \{\mathring{\mathbf{U}} \mid \mathbf{U} \in \mathcal{L}\} \quad (1)$$

Additionally, \mathcal{L} also represents a possibly infinite family of partial orders \mathcal{L}_{po} which is obtained by taking the transitive closure of each DAG in $\mathcal{L}_{\mathcal{G}}$.

$$\mathcal{L}_{po} = \{\mathbf{tc}(\mathring{\mathbf{U}}) \mid \mathbf{U} \in \mathcal{L}\} \quad (2)$$

A slice language $\mathcal{L} \subseteq \mathcal{L}(\vec{\Sigma}(c, T))$ is regular if it can be defined by a finite automaton \mathcal{A} over the slice alphabet $\vec{\Sigma}(c, T)$.

Definition 3.1 (Slice Automaton). Let T be a finite set of symbols and let $c \in \mathbb{N}$. A slice automaton over a slice alphabet $\vec{\Sigma}(c, T)$ is a finite automaton $\mathcal{A} = (Q, \Delta, q_0, F)$ where Q is a set of states, $q_0 \in Q$ is an initial state, $F \subseteq Q$ is a set of final states, and $\Delta \subseteq Q \times \vec{\Sigma}(c, T) \times Q$ is a transition relation such that for every $q, q', q'' \in Q$ and every $\mathbf{S} \in \vec{\Sigma}(c, T)$:

1. if $(q_0, \mathbf{S}, q) \in \Delta$ then \mathbf{S} is an initial slice,
2. if $(q, \mathbf{S}, q') \in \Delta$ and $q' \in F$, then \mathbf{S} is a final slice,
3. if $(q, \mathbf{S}, q') \in \Delta$ and $(q', \mathbf{S}', q'') \in \Delta$, then \mathbf{S} can be glued to \mathbf{S}' .

We denote by $\mathcal{L}(\mathcal{A})$ the slice language accepted by \mathcal{A} . We denote by $\mathcal{L}_{\mathcal{G}}(\mathcal{A})$ and $\mathcal{L}_{po}(\mathcal{A})$ respectively the set of DAGs derived from unit decompositions in $\mathcal{L}(\mathcal{A})$ and the set of partial orders obtained by taking the transitive closure of DAGs in $\mathcal{L}_{\mathcal{G}}(\mathcal{A})$.

4 Saturated and Transitively Reduced Slice Languages

Let H be a DAG whose vertices are labeled with elements from a finite set T . Then we let $\mathbf{ud}(H, c)$ denote the set of all unit decompositions \mathbf{U} in $\mathcal{L}(\vec{\Sigma}(c, T))$ for which $\mathbf{U} = H$. The set of all unit decompositions of H is defined as

$$\mathbf{ud}(H) = \bigcup_{c \geq 0} \mathbf{ud}(H, c) \quad (3)$$

We say that a slice language \mathcal{L} over $\vec{\Sigma}(c, T)$ is *saturated* if for every DAG $H \in \mathcal{L}_{\mathcal{G}}$ we have that $\mathbf{ud}(H) \subseteq \mathcal{L}$. Notice that if a slice language \mathcal{L} over $\vec{\Sigma}(c, T)$ is saturated, then for any $H \in \mathcal{L}_{\mathcal{G}}$ we have that $\mathbf{ud}(H, c) = \mathbf{ud}(H, c')$ for any $c' \geq c$. Let H be a DAG. An ordering $\omega = (v_1, v_2, \dots, v_n)$ of the vertices of H is a *topological ordering* if for any i, j with $1 \leq i < j \leq n$, there is no edge of H whose source is v_j and whose target is v_i . In other words, in a topological ordering, the target of an edge has always a greater position in the ordering than its source. Notice that if $\mathbf{U} = \mathbf{S}_1 \mathbf{S}_2 \dots \mathbf{S}_n$ is a unit decomposition of a DAG H , and if v_i is the center vertex of \mathbf{S}_i , then the ordering $\omega = (v_1, v_2, \dots, v_n)$ is always a topological ordering of H . We say that \mathbf{U} is compatible with ω . Conversely, given any topological ordering ω of H there exists at least one unit decomposition \mathbf{U} of H that is compatible with ω . We denote by $\mathbf{ud}(H, \omega)$ the set of all unit decompositions of H that are compatible with ω . Notice that $\mathbf{ud}(H) = \bigcup_{\omega} \mathbf{ud}(H, \omega)$ where ω ranges over all topological orderings of H . We say that a slice language is *vertically saturated* if for every $H \in \mathcal{L}_{\mathcal{G}}$ and every topological ordering ω of H , $\mathbf{ud}(H, \omega) \cap \mathcal{L} \neq \emptyset$ implies that $\mathbf{ud}(H, \omega) \subseteq \mathcal{L}$. Notice that a slice language may be vertically saturated without being saturated. In general, deriving from a slice automaton \mathcal{A} a slice automaton \mathcal{A}' such that $\mathcal{L}(\mathcal{A}')$ is saturated and such that $\mathcal{L}_{po}(\mathcal{A}') = \mathcal{L}_{po}(\mathcal{A})$ is an uncomputable problem [15]. However it is always possible to derive from \mathcal{A} a slice automaton \mathcal{A}'' such that $\mathcal{L}(\mathcal{A}'')$ is vertically saturated and such that $\mathcal{L}_{po}(\mathcal{A}'') = \mathcal{L}_{po}(\mathcal{A})$ [16]. We say that a slice automaton \mathcal{A} is saturated if $\mathcal{L}(\mathcal{A})$ is saturated. We say that \mathcal{A} is vertically saturated if $\mathcal{L}(\mathcal{A})$ is vertically saturated.

The transitive reduction of a DAG $H = (V, E, l)$ is the minimal subgraph $\mathbf{tr}(H)$ of H with the same transitive closure as H . In other words $\mathbf{tc}(\mathbf{tr}(H)) = \mathbf{tc}(H)$. We say

that a DAG H is transitively reduced if $H = \text{tr}(H)$. Alternatively, we say that a transitively reduced DAG is a Hasse diagram. We say that a slice language \mathcal{L} is transitively reduced if every DAG in \mathcal{L}_G is transitively reduced. The transitive reduction of a slice language \mathcal{L} is the unique slice language $\text{tr}(\mathcal{L})$ which is transitively reduced, vertically saturated and such that for each DAG $H \in \mathcal{L}_G$ and each topological ordering ω of H , $\text{ud}(H, \omega) \cap \mathcal{L} \neq \emptyset$ implies that $\text{ud}(\text{tr}(H), \omega) \cap \mathcal{L} \neq \emptyset$. Notice that by our definition of transitive reduction, if \mathcal{L} is saturated, then $\text{tr}(\mathcal{L})$ is also saturated. We say that a slice automaton \mathcal{A} is transitively reduced if $\mathcal{L}(\mathcal{A})$ is transitively reduced.

Lemma 4.1 (Transitive Reduction of Slice Languages [16]). *Let \mathcal{L} be a regular slice language represented by a finite automaton \mathcal{A} over $\vec{\Sigma}(c, T)$. Then there exists a finite automaton $\text{tr}(\mathcal{A})$ on $2^{O(c \log c)} \cdot |\mathcal{A}|$ states with $\mathcal{L}(\text{tr}(\mathcal{A})) = \text{tr}(\mathcal{L})$.*

Let T be a finite set of transitions. We denote by $\mathcal{P}(c, T)$ the set of all c -partial orders whose vertices are labeled with elements from T .

Lemma 4.2 ([16]). *For any finite set T and any $c \in \mathbb{N}$, one can construct a saturated transitively reduced slice automaton $\mathcal{A}(c, T)$ over $\vec{\Sigma}(c, T)$ such that $\mathcal{L}_{po}(\mathcal{A}(c, T)) = \mathcal{P}(c, T)$.*

Definition 4.3 (c -Complementation). *Let $\mathcal{P} \subseteq \mathcal{P}(c, T)$. Then we let $\overline{\mathcal{P}}^c = \mathcal{P}(c, T) \setminus \mathcal{P}$ be the c -complement of \mathcal{P} .*

The following lemma says that operations performed on transitively reduced saturated slice languages are reflected on the partial order languages they represent. Below $\mathcal{A} \cup \mathcal{A}'$, $\mathcal{A} \cap \mathcal{A}'$ and $\mathcal{A} \setminus \mathcal{A}'$ denote automata whose *slice language* (i.e. the syntactic language) is equal to $\mathcal{L}(\mathcal{A}) \cup \mathcal{L}(\mathcal{A}')$, $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{A}')$ and $\mathcal{L}(\mathcal{A}) \setminus \mathcal{L}(\mathcal{A}')$ respectively.

Lemma 4.4 (Properties of Saturated Slice Languages [16]). *Let \mathcal{A} and \mathcal{A}' be two transitively-reduced slice automata over $\vec{\Sigma}(c, T)$. Assume that \mathcal{A} is saturated.*

1. $\mathcal{L}_{po}(\mathcal{A} \cup \mathcal{A}') = \mathcal{L}_{po}(\mathcal{A}) \cup \mathcal{L}_{po}(\mathcal{A}')$
2. $\mathcal{L}_{po}(\mathcal{A} \cap \mathcal{A}') = \mathcal{L}_{po}(\mathcal{A}) \cap \mathcal{L}_{po}(\mathcal{A}')$
3. $\mathcal{L}_{po}(\mathcal{A}(c, T) \setminus \mathcal{A}) = \overline{\mathcal{L}_{po}(\mathcal{A})}^c$.
4. $\mathcal{L}_{po}(\mathcal{A}) \subseteq \mathcal{L}_{po}(\mathcal{A}')$ if and only if $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{A}')$.
5. $\mathcal{L}_{po}(\mathcal{A}) \cap \mathcal{L}_{po}(\mathcal{A}') = \emptyset$ if and only if $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{A}') = \emptyset$.
6. If \mathcal{A}' is saturated then $\mathcal{A} \cup \mathcal{A}'$ and $\mathcal{A} \cap \mathcal{A}'$ are also saturated.

Lemma 4.4 implies that union, intersection and c -complementation of partial order languages represented by transitively reduced saturated slice automata are computable, and inclusion and emptiness of intersection of these partial order languages are decidable. Theorem 4.5 establishes a close correspondence between the partial order behavior of bounded p/t -nets and regular slice languages.

Theorem 4.5 (*p/t-nets and Regular Slice Languages [15,16]*).

- i) **Expressibility:** Let $N = (P, T)$ be a b -bounded p/t -net and $sem \in \{ex, cau\}$. Then one can construct a saturated transitively reduced slice automaton $\mathcal{A}_{sem}(N, c)$ over $\vec{\Sigma}(c, T)$ such that $\mathcal{L}_{po}(\mathcal{A}_{sem}(N, c)) = \mathcal{P}_{sem}(N, c)$.
- ii) **Verification:** Let $N = (P, T)$ be a b -bounded p/t -net, \mathcal{A} be a slice automaton over $\vec{\Sigma}(c, T)$, and $sem \in \{ex, cau\}$.
 - (a) It is decidable whether $\mathcal{P}_{sem}(N, c) \cap \mathcal{L}_{po}(\mathcal{A}) = \emptyset$,
 - (b) It is decidable whether $\mathcal{L}_{po}(\mathcal{A}) \subseteq \mathcal{P}_{sem}(N, c)$,
 - (c) If \mathcal{A} is saturated then it is decidable whether $\mathcal{P}_{sem}(N, c) \subseteq \mathcal{L}_{po}(\mathcal{A})$.
- iii) **Synthesis:** Let \mathcal{A} be a slice automaton, $c, b, r \in \mathbb{N}$ and $sem \in \{ex, cau\}$. Then one may automatically determine whether there exists a (b, r) -bounded p/t -net N which is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A})$. In the case such a net exists, one may automatically construct it.

Observe that the synthesis result stated in Theorem 4.5.iii can be understood as the inverse of the expressibility result stated in Theorem 4.5.i. On the one hand, Theorem 4.5.i can be used to construct a slice automaton \mathcal{A} representing the c -partial order behavior of $N = (P, T)$. On the other hand, Theorem 4.5.iii can be used to recover from \mathcal{A} a p/t -net N' whose c -partial order behavior is equal to the c -partial-order behavior of N .

5 Monadic Second Order Logic of Graphs

The monadic second order logic of graphs MSO extends first order logic by allowing quantification over sets of vertices. The logic MSO_2 is an extension of MSO that also allows quantification over sets of edges. We refer to [12] for an extensive treatment of these logics. In this section we will use MSO_2 to describe properties of DAGs, while we will use MSO to describe properties of partial orders.

We will represent a partial order ℓ by a relational structure $\ell = (V, <, l)$ where V is a set of vertices, $< \subseteq V \times V$ is an ordering relation and $l \subseteq V \times T$ is a vertex labeling relation where T is a finite set of symbols (which should be regarded as the labels of transitions in a concurrent system). First order variables representing individual vertices will be taken from the set $\{x_1, x_2, \dots\}$ while second order variables representing sets of vertices will be taken from the set $\{X_1, X_2, \dots\}$. The set of MSO formulas is the smallest set of formulas containing:

- the atomic formulas $x_i \in X$, $x_i < x_j$, $l(x_i, a)$ for each $i, j \in \mathbb{N}$ with $i \neq j$ and each $a \in T$,
- the formulas $\varphi \wedge \psi$, $\varphi \vee \psi$, $\neg \varphi$, $\exists x_i. \varphi(x_i)$ and $\exists X_i. \varphi(X_i)$, where φ and ψ are MSO formulas.

An MSO sentence is a MSO formula φ without free variables. If φ is a sentence, and $\ell = (V, <, l)$ a partial order, then we denote by $\ell \models \varphi$ the fact that ℓ satisfies φ .

We will represent a general DAG G by a relational structure $G = (V, E, s, t, l)$ where V is a set of vertices, E a set of edges, $s, t \subseteq E \times V$ are respectively the source

and target relations, $l \subseteq V \times T$ is a vertex labeling relation, where T is a finite set of symbols. If e is an edge in E and v is a vertex in V then $s(e, v)$ is true if v is the source of e and $t(e, v)$ is true if v is the target of e . If $v \in V$ and $a \in T$ then $l(v, a)$ is true if v is labeled with a . First order variables representing individual vertices will be taken from the set $\{x_1, x_2, \dots\}$ and first order variables representing edges, from the set $\{y_1, y_2, \dots\}$. Second order variables representing sets of vertices will be taken from the set $\{X_1, X_2, \dots\}$ and second order variables representing sets of edges, from the set $\{Y_1, Y_2, \dots\}$. The set of MSO_2 formulas is the smallest set of formulas containing:

- the atomic formulas $x_i \in X_j, y_i \in Y_j, s(y_i, x_j), t(y_i, x_j), l(x_i, a)$ for each $i, j \in \mathbb{N}$ and $a \in T$,
- the formulas $\varphi \wedge \psi, \varphi \vee \psi, \neg \varphi, \exists x_i. \varphi(x_i)$ and $\exists X_i. \varphi(X_i), \exists y_i. \varphi(Y_i)$ and $\exists Y_i. \varphi(Y_i)$, where φ and ψ are MSO_2 formulas.

An MSO_2 sentence is a formula φ without free variables. If φ is a sentence, then we denote by $G \models \varphi$ the fact that G satisfies φ .

6 MSO Logic and Slice Languages

Lemma 6.1 below, which was proved in a more general context [17], states that the set of all unit decompositions \mathbf{U} whose graph $\mathring{\mathbf{U}}$ satisfy a given MSO_2 formula φ is a regular slice language.

Lemma 6.1 ([17]). *Given a MSO_2 formula φ , one can effectively construct a slice automaton $\mathcal{A}(c, T, \varphi)$ over $\vec{\Sigma}(c, T)$ such that*

$$\mathcal{L}(\mathcal{A}(c, T, \varphi)) = \{\mathbf{U} \in \mathcal{L}(\vec{\Sigma}(c, T)) \mid \mathring{\mathbf{U}} \models \varphi\}.$$

We say that a DAG $H = (V, E)$ can be covered by c paths if there exist simple paths p_1, \dots, p_c in H with $p_i = (V_i, E_i)$ such that $V = \cup_i V_i$ and $E = \cup_i E_i$. Proposition 6.2 below establishes a correspondence between c -coverable DAGs and their sets of unit decompositions.

Proposition 6.2. *Let H be a DAG. If H can be covered by c paths, then any unit decomposition of H has width at most c .*

We let $\gamma(c)$ be the MSO_2 sentence which is true on a DAG H whenever H can be covered by c paths. Then we have that $\mathcal{L}(\mathcal{A}(c, T, \varphi \wedge \gamma(c)))$ is the set of all unit decompositions in $\mathcal{L}(\mathcal{A}(c, T, \varphi))$ whose corresponding DAG can be covered by c -paths.

Lemma 6.3. *For any MSO_2 formula φ and any positive integer $c \in \mathbb{N}$, the slice automaton $\mathcal{A}(c, T, \varphi \wedge \gamma(c))$ is saturated.*

Recall that if H is a DAG, then $\mathbf{tr}(H)$ denotes the transitive reduction of H .

Proposition 6.4 (Partial Orders vs Hasse Diagrams). *For any MSO formula φ expressing a partial order property, there is an MSO_2 formula φ^{gr} expressing a property of DAGs such that for any partial order $\ell \in \mathcal{P}(c, T)$, $\ell \models \varphi$ if and only if $\mathbf{tr}(\ell) \models \varphi^{gr}$.*

Let $c \in \mathbb{N}$, T be a finite set, and φ be *MSO* formula. We denote by $\mathcal{P}(c, T, \varphi)$ the set of all c -partial orders satisfying φ whose vertices are labeled with elements from T . We denote by ρ be the *MSO*₂ formula which is true on a DAG H whenever H is transitively reduced, i.e., whenever $H = \text{tr}(H)$.

Lemma 6.5. *Let φ be a *MSO* formula expressing a partial order property, and φ^{gr} be the *MSO*₂ formula of Proposition 6.4. Then $\mathcal{A}(c, T, \varphi^{gr} \wedge \rho \wedge \gamma(c))$ is a saturated transitively reduced slice automaton and $\mathcal{P}(c, T, \varphi) = \mathcal{L}_{po}(\mathcal{A}(c, T, \varphi^{gr} \wedge \rho \wedge \gamma(c)))$.*

Lemma 6.6 (Verifying Regular Slice Languages). *Let φ be a *MSO* formula, and let \mathcal{A} be a transitively reduced saturated slice automaton over $\vec{\Sigma}(c, T)$.*

- i) *One may effectively verify whether $\mathcal{L}_{po}(\mathcal{A}) \cap \mathcal{P}(c, T, \varphi) = \emptyset$.*
- ii) *One may effectively verify whether $\mathcal{L}_{po}(\mathcal{A}) \subseteq \mathcal{P}(c, T, \varphi)$.*
- iii) *One may effectively verify whether $\mathcal{P}(c, T, \varphi) \subseteq \mathcal{L}_{po}(\mathcal{A})$.*

7 Proofs of our Main Results

Finally, we are in a position to prove our main results. First we will state a lemma that we call the separation lemma.

Lemma 7.1 (Separation Lemma). *Let \mathcal{A} and \mathcal{A}' be two slice automata over $\vec{\Sigma}(c, T)$. And suppose that \mathcal{A}' is saturated. Then one can decide whether there exists a (b, r) -bounded p/t -net N such that*

- i) *N is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A})$,*
- ii) *$\mathcal{P}(N, c) \cap \mathcal{L}_{po}(\mathcal{A}') = \emptyset$.*

In the case such a net N exists one can automatically construct it.

Proof. First we apply Theorem 4.5.iii to determine if there exists a (b, r) -bounded p/t -net N that is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A})$. In case such a net exists we construct it. Now using Theorem 4.5.i we construct a slice automaton \mathcal{A}'' show causal/execution behavior is precisely that of N . Finally, since \mathcal{A}' is transitively reduced and saturated we can use Lemma 4.4 to test whether $\mathcal{L}_{po}(\mathcal{A}') \cap \mathcal{L}_{po}(\mathcal{A}'')$. Notice that by the minimality of $\mathcal{L}_{po}(N)$ if this intersection is not empty, then the problem has no solution. \square

Proof of Theorem 1.1 Let $N = (P, T)$ be a b -bounded p/t -net and let $sem \in \{ex, cau\}$. By Theorem 4.5 we can construct a saturated, transitively reduced slice automaton $\mathcal{A}_{sem}(N, c)$ such that $\mathcal{L}_{po}(\mathcal{A}_{sem}(N, c)) = \mathcal{P}_{sem}(N, c)$. Now by Lemma 6.6 we can effectively determine whether $\mathcal{L}_{po}(\mathcal{A}_{sem}(N, c)) \cap \mathcal{P}(c, T, \varphi) = \emptyset$, whether $\mathcal{L}_{po}(\mathcal{A}_{sem}(N, c)) \subseteq \mathcal{P}(c, T, \varphi)$ or whether $\mathcal{P}(c, T, \varphi) \subseteq \mathcal{L}_{po}(\mathcal{A}_{sem}(N, c))$. \square

Proof of Theorem 1.2 Let φ be a *MSO* formula. By Lemma 6.5 one can construct a saturated, transitively reduced slice automaton $\mathcal{A} = \mathcal{A}(c, T, \varphi \wedge \rho \wedge \gamma(c))$ over $\vec{\Sigma}(c, T)$ such that $\mathcal{L}_{po}(\mathcal{A}) = \mathcal{P}(c, T, \varphi)$. By Theorem 4.5 one may automatically determine whether there exists a (b, r) -bounded p/t -net N which is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A})$, and in the case such a net exists, one may automatically construct it. \square

Proof of Theorem 1.3: By Lemma 6.5 one can construct a saturated, transitively reduced slice automaton $\mathcal{A} = \mathcal{A}(c, T, \varphi^{gr} \wedge \rho \wedge \gamma(c))$ such that $\mathcal{L}_{po}(\mathcal{A}) = \mathcal{P}(c, T, \varphi)$. By Theorem 4.5.i, one can construct a saturated, transitively reduced slice automaton \mathcal{A}' such that $\mathcal{L}_{po}(\mathcal{A}') = \mathcal{P}_{sem}(N, c)$. Since both \mathcal{A} and \mathcal{A}' are saturated and transitively reduced, by Lemma 4.4, we have that the slice automaton $\mathcal{A} \cap \mathcal{A}'$ is saturated and transitively reduced. Additionally $\mathcal{L}_{po}(\mathcal{A} \cap \mathcal{A}') = \mathcal{L}_{po}(\mathcal{A}) \cap \mathcal{L}_{po}(\mathcal{A}') = \mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c)$. Additionally, by Lemma 4.4 we can construct a transitively reduced and saturated slice automaton $\overline{\mathcal{A}'}^c$ such that $\mathcal{L}(\overline{\mathcal{A}'}^c) = \mathcal{P}(c, T) \setminus po(N, c)$. Thus as a last step we may apply Lemma 7.1 to determine whether there exists a (b, r) -bounded p/t -net N that is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A} \cap \mathcal{A}')$ and such that $\mathcal{P}(N, c) \cap \mathcal{L}_{po}(\overline{\mathcal{A}'}^c) = \emptyset$, and in the case that such a net exists, we can effectively construct it. \square

Proof of Theorem 1.4: By Lemma 6.5 one can construct a saturated, transitively reduced slice automata $\mathcal{A}_\varphi = \mathcal{A}(c, T, \varphi^{gr} \wedge \rho \wedge \gamma(c))$ and $\mathcal{A}_\psi = \mathcal{A}(c, T, \psi^{gr} \wedge \rho \wedge \gamma(c))$ such that $\mathcal{L}_{po}(\mathcal{A}_\varphi) = \mathcal{P}(c, T, \varphi)$ and $\mathcal{L}_{po}(\mathcal{A}_\psi) = \mathcal{P}(c, T, \psi)$ respectively. By Theorem 4.5.i, one can construct a saturated, transitively reduced slice automaton \mathcal{A}' such that $\mathcal{L}_{po}(\mathcal{A}') = \mathcal{P}_{sem}(N, c)$. Since both \mathcal{A} and \mathcal{A}' are saturated and transitively reduced, by Lemma 4.4, we have that the slice automaton $\mathcal{A}_\varphi \cap \mathcal{A}'$ is saturated and transitively reduced. Additionally $\mathcal{L}_{po}(\mathcal{A}_\varphi \cap \mathcal{A}') = \mathcal{L}_{po}(\mathcal{A}_\varphi) \cap \mathcal{L}_{po}(\mathcal{A}') = \mathcal{P}(c, T, \varphi) \cap \mathcal{P}_{sem}(N, c)$. Thus as a last step we may apply Lemma 7.1 to determine whether there exists a (b, r) -bounded p/t -net N that is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A}_\varphi \cap \mathcal{A}')$ and such that $\mathcal{P}(N, c) \cap \mathcal{L}_{po}(\mathcal{A}_\psi) = \emptyset$, and in the case that such a net exists, we can effectively construct it. \square

Proof of Theorem 1.5: Let φ^{yes} and φ^{no} be two MSO formulas specifying respectively a set of good partial order behaviors and a set of bad partial order behaviors. By Lemma 6.5 we can construct saturated, transitively reduced slice automata

$$\mathcal{A}^{yes} = \mathcal{A}(c, T, [\varphi^{yes}]^{gr} \wedge \rho \wedge \gamma(c)) \quad \text{and} \quad \mathcal{A}^{no} = \mathcal{A}(c, T, [\varphi^{no}]^{gr} \wedge \rho \wedge \gamma(c))$$

such that $\mathcal{L}_{po}(\mathcal{A}^{yes}) = \mathcal{P}(c, T, \varphi^{yes})$ and $\mathcal{L}_{po}(\mathcal{A}^{no}) = \mathcal{P}(c, T, \varphi^{no})$. Now by Theorem 4.5.iii we can synthesize a (b, r) -bounded p/t -net N that is c -sem-minimal with respect to $\mathcal{L}_{po}(\mathcal{A}^{yes})$. Since \mathcal{A}^{no} is saturated we can apply Lemma 7.1 to determine whether there exists a (b, r) -bounded p/t -net N that is c -sem-minimal for $\mathcal{L}_{po}(\mathcal{A}^{yes})$ and such that $\mathcal{P}(N, c) \cap \mathcal{L}_{po}(\mathcal{A}^{no}) = \emptyset$. In the case such a net exists we can effectively construct it. \square

8 Conclusion

In this work we have shown that both model checking of the c -partial-order behavior of bounded p/t -nets and the synthesis of bounded p/t -nets from MSO definable sets of c -partial-orders are computationally feasible. By combining these two results, we introduced the *semantically safest subsystem* problem as a new primitive for the study of automated correction of computational systems. Additionally we were able to lift the theory of automatic program repair developed in [37] to the realm of bounded p/t -nets and to develop a methodology of synthesis by contracts that is suitable for the partial order theory of concurrency.

References

1. J. Alglave, D. Kroening, and M. Tautschnig. Partial orders for efficient bounded model checking of concurrent software. In *CAV 2013*, volume 8044 of *LNCS*, pages 141–157. Springer, 2013.
2. F. Avellaneda and R. Morin. Checking partial-order properties of vector addition systems with states. In *ACSD 2013*, pages 100–109. IEEE, 2013.
3. E. Badouel and P. Darondeau. On the synthesis of general Petri nets. Technical Report PI-1061, IRISA, 1996.
4. E. Badouel and P. Darondeau. Theory of regions. In *Lectures on Petri Nets I: Basic Models*, volume 1491 of *LNCS*, pages 529–586. Springer, 1998.
5. R. Bergenthum, J. Desel, R. Lorenz, and S. Mauser. Synthesis of Petri nets from finite partial languages. *Fundamenta Informaticae*, 88(4):437–468, 2008.
6. R. Bergenthum, J. Desel, R. Lorenz, and S. Mauser. Synthesis of Petri nets from infinite partial languages. In *ACSD 2008*, pages 170–179. IEEE, 2008.
7. J. R. Büchi. Weak second order arithmetic and finite automata. *Z. Math. Logik Grundl. Math.*, 6:66–92, 1960.
8. P. Cerný, K. Chatterjee, T. A. Henzinger, A. Radhakrishna, and R. Singh. Quantitative synthesis for concurrent programs. In *CAV 2011*, volume 6806 of *LNCS*, pages 243–259. Springer, 2011.
9. P. Cerný, T. A. Henzinger, A. Radhakrishna, L. Ryzhyk, and T. Tarrach. Efficient synthesis for concurrency by semantics-preserving transformations. In *CAV 2013*, volume 8044 of *LNCS*, pages 951–967. Springer, 2013.
10. E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
11. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. MIT press, 1999.
12. B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach*, volume 138. Cambridge University Press, 2012.
13. P. Darondeau. Deriving unbounded Petri nets from formal languages. *LNCS*, 1466:533–548, 1998.
14. P. Darondeau. Region based synthesis of P/T-nets and its potential applications. In *ICATPN 2000*, volume 1825 of *LNCS*, pages 16–23. Springer, 2000.
15. M. de Oliveira Oliveira. Hasse diagram generators and Petri nets. *Fundamenta Informaticae*, 105(3):263–289, 2010.
16. M. de Oliveira Oliveira. Canonizable partial order generators. In *LATA 2012*, volume 7183 of *LNCS*, pages 445–457. Springer, 2012.
17. M. de Oliveira Oliveira. Subgraphs satisfying mso properties on \mathbb{Z} -topologically orderable digraphs. In *IPEC 2013*, volume 8246 of *LNCS*, pages 123–136. Springer, 2013.
18. E. A. Emerson and E. M. Clarke. Using branching time temporal logic to synthesize synchronization skeletons. *Science of Computer programming*, 2(3):241–266, 1982.
19. M.-P. Flé and G. Roucairol. On serializability of iterated transactions. In *PODC 1982*, pages 194–200. ACM, 1982.
20. H. Gaifman and V. R. Pratt. Partial order models of concurrency and the computation of functions. In *LICS 1987*, pages 72–85, 1987.
21. J. L. Gischer. The equational theory of pomsets. *Theoretical Computer Science*, 61:199–224, 1988.
22. U. Goltz and W. Reisig. Processes of place/transition-nets. In *ICALP 1983*, volume 154, pages 264–277, 1983.

23. A. Griesmayer, R. Bloem, and B. Cook. Repair of boolean programs with an application to C. In *CAV 2006*, volume 4144 of *LNCS*, pages 358–371. Springer, 2006.
24. B. Jobstmann, A. Griesmayer, and R. Bloem. Program repair as a game. In *CAV 2005*, volume 3576 of *LNCS*, pages 226–238. Springer, 2005.
25. G. Juhás, R. Lorenz, and J. Desel. Can I execute my scenario in your net? In *ICATPN 2005*, volume 3536 of *LNCS*, pages 289–308, 2005.
26. O. Kupferman, Y. Lustig, M. Y. Vardi, and M. Yannakakis. Temporal synthesis for bounded systems and environments. In *STACS 2011*, volume 9 of *LIPIcs*, pages 615–626, 2011.
27. Madhusudan. Reasoning about sequential and branching behaviours of message sequence graphs. In *Proc. of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, volume 2076 of *LNCS*, pages 809–820, 2001.
28. Z. Manna and P. Wolper. Synthesis of communicating processes from temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 6(1):68–93, 1984.
29. T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.
30. C. A. Petri. Fundamentals of a theory of asynchronous information flow. In *Proceedings of IFIP Congress 62*, pages 166–168, Munchen, 1962.
31. A. Pnueli. The temporal semantics of concurrent programs. *Theoretical computer science*, 13(1):45–60, 1981.
32. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL 1989*, pages 179–190. ACM, 1989.
33. J.-P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *International Symposium on Programming*, pages 337–351. Springer, 1982.
34. R. Samanta, J. V. Deshmukh, and E. A. Emerson. Automatic generation of local repairs for boolean programs. In *FMCAD 2008*, pages 1–10. IEEE, 2008.
35. R. Van Glabbeek and U. Goltz. Refinement of actions and equivalence notions for concurrent systems. *Acta Informatica*, 37(4-5):229–327, 2001.
36. W. Vogler. *Modular Construction and Partial Order Semantics of Petri Nets*, volume 5606 of *LNCS*. 2009.
37. C. von Essen and B. Jobstmann. Program repair without regret. In *CAV 2013*, volume 8044 of *LNCS*, pages 896–911. Springer, 2013.
38. R. Zurawski and M. Zhou. Petri nets and industrial applications: A tutorial. *Industrial Electronics, IEEE Transactions on*, 41(6):567–583, 1994.

A Proofs of Auxiliary Results

Proof of Proposition 6.2: Let H be the union of c paths p_1, \dots, p_c and let $U = S_1 S_2 \dots S_n$ be a unit decomposition of H . Let v_i be the center vertex of S_i . Then the ordering $\omega = (v_1, \dots, v_n)$ is a topological ordering of H . This implies that for any $i \in \{1, \dots, n\}$, and any $j \in \{1, \dots, c\}$ there exists at most one edge from p_j whose source is in $\{v_1, \dots, v_i\}$ and whose target is in $\{v_{i+1}, \dots, v_n\}$. Thus there exists at most c edges in $p_1 \cup \dots \cup p_c$ with whose source is in $\{v_1, \dots, v_i\}$ and whose target is in $\{v_{i+1}, \dots, v_n\}$. This implies that $|E(\{v_1, \dots, v_i\}, \{v_{i+1}, \dots, v_n\})| \leq c$ for each $i \in \{1, \dots, n\}$. and thus ω has cut-width at most c with respect to H since the $w(U)$ is equal to the cut-width of ω we have that that $w(U) \leq c$. \square

Proof of Lemma 6.3: By Lemma 6.1, $\mathcal{L}(c, T, \varphi \wedge \gamma(c))$ is a regular slice language over $\vec{\Sigma}(c, T)$. Thus we just need to show that $\mathcal{L}(c, T, \varphi \wedge \gamma(c))$ is saturated. A unit decomposition \mathbf{U} belongs to $\mathcal{L}(c, T, \varphi \wedge \gamma(c))$ if and only if \mathbf{U} satisfies the following three properties: $\mathbf{U} \in \mathcal{L}(c, T)$, \mathbf{U} can be covered by c paths and $\mathbf{U} \models \varphi$. Since \mathbf{U} can be covered by c paths, it follows from Proposition 6.2 that $\mathbf{ud}(\mathbf{U}) \subseteq \mathcal{L}(\vec{\Sigma}(c, T))$. Now let \mathbf{U}' be an arbitrary unit decomposition in $\mathbf{ud}(\mathbf{U})$. Since $\mathbf{U}' = \mathbf{U}$, we have that \mathbf{U}' is the union of c paths and satisfies φ . Thus $\mathbf{U}' \in \mathcal{L}(c, T, \varphi \wedge \gamma(c))$. Since \mathbf{U}' was taken to be an arbitrary unit decomposition in $\mathbf{ud}(\mathbf{U})$, we have that $\mathcal{L}(c, T, \varphi \wedge \gamma(c))$ is saturated. \square

Proof of Proposition 6.4: Let $\ell = (V, <, l)$ be a partial order and $\mathbf{tr}(\ell) = (V, E, l)$ be the transitive reduction of ℓ . Then for any two vertices $v, v' \in V$, we have that $v < v'$ if and only if there is a path $v = v_1 e_1 v_2 \dots e_{n-1} v_n = v'$ from v to v' in $\mathbf{tr}(\ell)$. Now let $\text{path}(x_1, X, Y, x_2)$ be a MSO_2 formula which is true in a DAG H whenever there is a path starting at x_1 , finishing at x_2 , with internal vertices X and internal edges Y . For a MSO formula φ , let φ^{gr} be the MSO_2 formula which is obtained from φ by replacing each occurrence of the atomic formula $x_1 < x_2$ in φ by the formula $\exists X \exists Y \text{path}(x_1, X, Y, x_2)$. Now have that $\ell \models \varphi$ if and only if $\mathbf{tr}(\ell) \models \varphi^{gr}$. \square

Proof of Lemma 6.5: Let ρ be the MSO_2 formula which is true in a DAG H whenever H is transitively reduced. Let φ^{gr} be the formula obtained from φ as in Proposition 6.4. Then we have that a DAG H satisfies $\gamma(c) \wedge \rho \wedge \varphi^{gr}$ if and only if H can be covered by c paths, H is transitively reduced and if the partial order $\mathbf{tc}(H)$ induced by H satisfies φ . By Lemma 6.3, the slice language $\mathcal{L}(c, T, \varphi \wedge \rho \wedge \gamma(c))$ is saturated, regular, and consists precisely of the unit decompositions yielding a graph satisfying $\varphi \wedge \rho \wedge \psi(c)$. Thus we just need to set $\mathcal{A}(c, T, \varphi)$ as the minimal deterministic finite automaton generating $\mathcal{L}(c, T, \varphi \wedge \rho \wedge \gamma(c))$. \square

Proof of Lemma 6.6: Let $\mathcal{A}' = \mathcal{A}(c, T, \varphi^{gr} \wedge \rho \wedge \gamma(c))$. By Lemma 6.5 \mathcal{A}' is saturated, transitively reduced and $\mathcal{L}(\mathcal{A}') = \mathcal{P}(c, T, \varphi)$. Since \mathcal{A} is also transitively reduced and saturated, by Lemma 4.4, $\mathcal{L}_{po}(\mathcal{A}) \cap \mathcal{L}_{po}(\mathcal{A}') = \emptyset$ if and only if $\mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{A}') = \emptyset$, $\mathcal{L}_{po}(\mathcal{A}) \subseteq \mathcal{L}_{po}(\mathcal{A}')$ if and only if $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{A}')$ and $\mathcal{L}_{po}(\mathcal{A}') \subseteq \mathcal{L}_{po}(\mathcal{A})$ if and only if $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{A}')$. Thus we have reduced emptiness of intersection and inclusion of the partial order languages represented by \mathcal{A} and \mathcal{A}' to the emptiness of intersection and inclusion of the regular slice languages accepted by \mathcal{A} and \mathcal{A}' . \square